Accelerating cross-validation with total variation and its application to super-resolution imaging

Tomoyuki Obuchi[1],* Shiro Ikeda[2], Kazunori Akiyama[3,4], and Yoshiyuki Kabashima[1]
[1]*Department of Mathematical and Computing Science, Tokyo Institute of Technology, Yokohama 226-8502, Japan*
[2]*The Institute of Statistical Mathematics, Tachikawa, Tokyo, 190-8562, Japan*
[3]*Haystack Observatory, Massachusetts Institute of Technology, Westford, MA 01886, USA*
[4]*National Astronomy Observatory of Japan, Osawa 2-21-1, Mitaka, Tokyo 181-8588, Japan*
(Dated: November 23, 2016)

We develop an approximation formula for the cross-validation error (CVE) of a sparse linear regression penalized by ℓ_1-norm and total variation terms, which is based on a perturbative expansion utilizing the largeness of both the data dimensionality and the model. The developed formula allows us to reduce the necessary computational cost of the CVE evaluation significantly. The practicality of the formula is tested through application to simulated black-hole image reconstruction on the event-horizon scale with super resolution. The results demonstrate that our approximation reproduces the CVE values obtained via literally conducted cross-validation with reasonably good precision.

arXiv:1611.07197v1 [stat.ME] 22 Nov 2016

Introduction At present, in many practical situations of science and technology, large high-dimensional observational datasets are created and accumulated on a continuous basis. An essential difficulty in treating such high-dimensional data is the extraction of meaningful information. Sparse modeling [1, 2] is a promising framework for overcoming this difficulty, which has recently been utilized in many disciplines [3, 4]. In this framework, a statistical or machine-learning model with a large number of parameters (explanatory variables) is fitted to the data, in conjunction with a certain sparsity-inducing penalty. This penalty should be appropriately chosen with consideration of the processed data. One representative penalty is the ℓ_1 regularization, which retains certain preferred properties, such as the statistical model convexity [5, 6]. A similar penalty attracting much attention is the so-called "total variation (TV)" [7–9]. The TV yields "continuity" of the neighboring variables, which is suitable for the processing of certain datasets expected to have such continuity, such as natural images [4, 7–9].

Another common difficulty associated with the use of statistical models is model selection. In the context of image processing using the ℓ_1 and TV regularizations, this difficulty appears during the selection of appropriate regularization weights. A practical framework to select these weights, which is applicable to general situations, is cross-validation (CV). CV provides an estimator of the statistical-model generalization error, i.e., the CV error (CVE), using the data under control, and the minimum CVE obtained when sweeping the weights yields the optimal weight values. This versatile framework is, however, computationally demanding for large datasets/models. This problem frequently becomes a bottleneck affecting model selection. Thus, reducing the CVE computational cost could have a significant impact on a broad range of applications in various disciplines.

Considering these circumstances, in this paper, we provide a CVE approximation formula for a linear regression penalized by the ℓ_1 and TV terms, to efficiently reduce the computational cost. The formula for the case penalized by the ℓ_1 term alone has already been proposed in [10], and the formula presented herein is a generalization of it. Below, we show the formula derivation and perform a demonstration in the context of super-resolution imaging. The processed images employed in this study are reconstructed from simulated observations of black holes on the event-horizon scale for the Event Horizon Telescope (EHT, see [11–13]) full array. Note that our formula will be applied to actual EHT observations to be conducted in April 2017.

Problem setting Let us suppose that our measurement is a linear process, and denote the measurement result as $\boldsymbol{y} \in \mathbb{R}^M$ and the measurement matrix as $A = \{A_{\mu i}\}_{\mu=1,\cdots,M;\, i=1,\cdots N} \in \mathbb{R}^{M\times N}$. The explanatory variables, corresponding to the images that will be examined in the later demonstration, are denoted by $\boldsymbol{x} \in \mathbb{R}^N$. The quality of the fit to the data is described by the residual sum of squares (RSS), i.e., $\mathcal{E}(\boldsymbol{x}|\boldsymbol{y}, A) = \frac{1}{2}||\boldsymbol{y} - A\boldsymbol{x}||_2^2$. In addition, we consider the following penalty consisting of ℓ_1 and TV terms:

$$R(\boldsymbol{x}; \lambda_{\ell_1}, \lambda_T) = \lambda_{\ell_1}||\boldsymbol{x}||_1 + \lambda_T T(\boldsymbol{x}), \tag{1}$$

where the $T(\boldsymbol{x})$ term corresponds to the TV and is expressed as

$$T(\boldsymbol{x}) = \sum_i \sqrt{\sum_{j\in\partial i}(x_j - x_i)^2} \equiv \sum_i t_i, \tag{2}$$

and ∂i denotes the neighboring variables of the ith variable. There is some variation in the definition of "neighbors"; here, we follow the standard approach [7–9]. That is, $\boldsymbol{x}$ is assumed to be a two-dimensional image and the neighbors of the ith pixel correspond to the right and down pixels. However, the bottom row (the rightmost column) of the image is exceptional, as the neighbor of

each pixel in that row (column) corresponds to the right (down) pixel only. Note that the developed approximation formula presented below is independent of this specific choice of neighbors and can be applied to general cases.

For this setup, we consider the following linear regression problem with the penalty given in Eq. (1)

$$\hat{\boldsymbol{x}}(\lambda_{\ell_1},\lambda_T)=\arg\min_{\boldsymbol{x}}\left\{\mathcal{E}(\boldsymbol{x}|\boldsymbol{y},A)+R(\boldsymbol{x};\lambda_{\ell_1},\lambda_T)\right\}, \quad (3)$$

where $\arg\min_u\{f(u)\}$ generally represents the argument that minimizes an arbitrary function $f(u)$. Further, we consider the leave-one-out (LOO) CV of Eq. (3) in the form

$$\hat{\boldsymbol{x}}^{\backslash\mu}(\lambda_{\ell_1},\lambda_T)=\arg\min_{\boldsymbol{x}}\left\{\frac{1}{2}\sum_{\nu(\neq\mu)}\left(y_\nu-A_{\nu i}x_i\right)^2+R\right\} \quad (4)$$

Note that the system without the μth row of $\boldsymbol{y}$ and A is referred to as the "μth LOO system" hereafter. In this procedure, the CVE, i.e., the generalization error estimator, is

$$\mathcal{E}_{\rm LOO}(\lambda_{\ell_1},\lambda_T)=\frac{1}{2}\sum_{\mu=1}^{M}(y_\mu-\boldsymbol{a}_\mu^{\rm T}\hat{\boldsymbol{x}}^{\backslash\mu}(\lambda_{\ell_1},\lambda_T))^2, \quad (5)$$

where $\boldsymbol{a}_\mu^\top=(A_{\mu1},\cdots,A_{\mu N})$ is the μth row vector of A. We term this simply the "LOO error (LOOE)."

Approximation formula for softened system When M is sufficiently large, i.e., the number of observations is large enourgh, the difference between the LOO solution $\hat{\boldsymbol{x}}^{\backslash\mu}$ and the full solution $\hat{\boldsymbol{x}}$ is expected to be small. This intuition naturally motivates us to conduct a perturbation connecting these two solutions. To conduct this perturbation, we "soften" the penalty by introducing a small cutoff $\delta(>0)$ in the TV, having the form

$$T\to T^\delta=\sum_i\sqrt{\sum_{j\in\partial i}(x_j-x_i)^2+\delta^2}\equiv\sum_i t_i^\delta. \quad (6)$$

An approximation formula in the presence of ℓ_1 regularization with smooth cost functions has already been proposed in [10]. We employ that formula here and take the limit $\delta\to0$.

To state the approximation formula, we begin by defining "active" and "killed" variables. Owing to the ℓ_1 term, some variables are set to zero in $\hat{\boldsymbol{x}}$; we refer to these variables as "killed variables." The remaining finite variables are termed "active variables." We denote the index sets of the active and killed variables by S_A and S_K, respectively. The active (killed) components of a vector $\boldsymbol{x}$ are formally expressed as $\boldsymbol{x}_{S_A}(\boldsymbol{x}_{S_K})$. For any matrix X, we use double subscripts in the same manner. For example, a submatrix having row and column components of S_A and S_K, respectively, is denoted $X_{S_AS_K}$.

The approximation formula can be derived through the following two steps. Note that, in this derivation, a crucial assumption is that the sets of active and killed variables are common among the full and LOO systems. This assumption may not hold exactly in practice, but the resultant formula is asymptotically exact in the large-N limit [10].

The first step is to compute the values of the active variables and their response to small perturbation. The active variables are determined by the extremization condition of the softened cost function with respect to the active variables, such that

$$\frac{\partial\left(\mathcal{E}(\boldsymbol{x}|\boldsymbol{y}^{\backslash\mu},A^{\backslash\mu})+R^\delta(\boldsymbol{x};\lambda_{\ell_1},\lambda_T)\right)}{\partial\boldsymbol{x}_{S_A}}=\boldsymbol{0}\Rightarrow(\hat{\boldsymbol{x}}^{\delta\backslash\mu})_{S_A} \quad (7)$$

The focus here is the response of this solution when a small perturbation $-\boldsymbol{h}\cdot\boldsymbol{x}$ is incorporated into the cost function. A simple computation demonstrates that the active–active components of the response function, $\left(\chi^{\delta\backslash\mu}\right)_{S_AS_A}=\frac{\partial}{\partial\boldsymbol{h}_{S_A}}(\hat{\boldsymbol{x}}^{\delta\backslash\mu})_{S_A}\Big|_{\boldsymbol{h}=0}$, are equivalent to the inverse of the cost-function Hessian

$$\begin{aligned}H^{\delta\backslash\mu}&=\partial_{\boldsymbol{x}}^2\left(\mathcal{E}(\boldsymbol{x}|\boldsymbol{y}^{\backslash\mu},A^{\backslash\mu})+R^\delta(\boldsymbol{x};\lambda_{\ell_1},\lambda_T)\right),\\&=G^{\backslash\mu}+\partial_{\boldsymbol{x}}^2R^\delta(\boldsymbol{x};\lambda_{\ell_1},\lambda_T),\end{aligned} \quad (8)$$

where $\partial_{\boldsymbol{x}}^2$ denotes the Hessian operator $\partial_{\boldsymbol{x}}^2\equiv(\frac{\partial^2}{\partial x_i\partial x_j})$ and $G^{\backslash\mu}$ is the Gram matrix of $A^{\backslash\mu}$, i.e., $G^{\backslash\mu}\equiv\left(A^{\backslash\mu}\right)^\top A^{\backslash\mu}$. The other components of the response function are identically zero, from the stability assumption of S_K and because the killed variables are zero, with $\hat{\boldsymbol{x}}_{S_K}=\hat{\boldsymbol{x}}_{S_K}^{\backslash\mu}=\boldsymbol{0}$.

In the second step, we connect the full solution to the LOO solution, through the above perturbation with an appropriate $\boldsymbol{h}$. To specify the perturbation, we assume that the difference $\boldsymbol{d}^\delta=\hat{\boldsymbol{x}}^\delta-\hat{\boldsymbol{x}}^{\delta\backslash\mu}$ is small and expand the RSS of the full system with respect to $\boldsymbol{d}^\delta$ as follows:

$$\mathcal{E}(\hat{\boldsymbol{x}}^{\delta\backslash\mu}|\boldsymbol{y},A)\approx\mathcal{E}(\hat{\boldsymbol{x}}^\delta|\boldsymbol{y},A)-\sum_{\mu=1}^{M}(y_\mu-\boldsymbol{a}_\mu^\top\hat{\boldsymbol{x}}^\delta)\boldsymbol{a}_\mu^\top\boldsymbol{d}^\delta. \quad (9)$$

This equation implies that the perturbation between the full and LOO systems can be expressed as $\boldsymbol{h}_\mu=(y_\mu-\boldsymbol{a}_\mu^\top\hat{\boldsymbol{x}}^\delta)\boldsymbol{a}_\mu$. Hence, we obtain

$$\hat{\boldsymbol{x}}^\delta\approx\hat{\boldsymbol{x}}^{\delta\backslash\mu}+\chi^{\delta\backslash\mu}\boldsymbol{h}_\mu=\hat{\boldsymbol{x}}^{\delta\backslash\mu}+(y_\mu-\boldsymbol{a}_\mu^\top\hat{\boldsymbol{x}}^\delta)\chi^{\delta\backslash\mu}\boldsymbol{a}_\mu. \quad (10)$$

The Hessian of the full system has a simple relationship with the LOO Hessian, such that

$$H^\delta\approx H^{\delta\backslash\mu}+\left(\boldsymbol{a}_\mu\boldsymbol{a}_\mu^\top\right), \quad (11)$$

Inserting Eqs. (10) and (11) in conjunction with $\chi^\delta_{S_AS_A}=(H^\delta_{S_AS_A})^{-1}$ into Eq. (5) and using the Sherman-Morrison formula for matrix inversion, we find

$$\mathcal{E}_{\rm LOO}(\lambda_{\ell_1},\lambda_T)\approx\frac{1}{2}\sum_{\mu=1}^{M}\frac{(y_\mu-\boldsymbol{a}_\mu^\top\hat{\boldsymbol{x}}^\delta)^2}{\left(1-\boldsymbol{a}_{\mu S_A}^\top(\chi^\delta)_{S_AS_A}\boldsymbol{a}_{\mu S_A}\right)^2}. \quad (12)$$

According to Eq. (12), we can compute the LOOE only from the full solution $\hat{\boldsymbol{x}}^\delta$, without actually performing CV, which facilitates considerable reduction of the computational cost.

Handling a singularity Let us generalize Eq. (12) to the limit $\delta \to 0$, where the penalty contains another singular term in addition to the ℓ_1 term. This TV singularity tends to "lock" some of the neighboring variables, i.e., $x_j = x_i$ $(\forall j \in \partial i)$, which corresponds to $t_i = 0$ in Eq. (2). If two different vanishing TV terms, t_i and t_j, share a common variable x_r, all the variables in those TV terms take the same value $x_k = x_r$ $(\forall k \in (\{i\} \cup \partial i \cup \{j\} \cup \partial j))$. In this manner, the active variables are separated into several "locked" clusters, with all the variables inside a cluster having an identical value. This implies that the variable response to a perturbation, $\chi = \lim_{\delta\to 0} \chi^\delta$, should have the same value for all variables in a cluster and may, therefore, be merged. Below, we demonstrate this behavior for the $\delta \to 0$ limit. For the derivation, we assume that the clusters are common to both the full and LOO systems, similar to the assumption for S_A and S_K. For convenience, we index the clusters by $\alpha, \beta \in C$ and denote the number of clusters by $|C|$; the index set of variables in a cluster α is represented by S_α and the total set of indices in all clusters is denoted by $S_C \equiv \cup_\alpha S_\alpha$. Hereafter, we concentrate on the active variable space only and omit the killed variable space. The complement set of S_C, i.e., the set of isolated variables that do not belong to any cluster, is denoted by S_I and, thus, $S_A = S_I \cup S_C$.

Two crucial observations for the derivation are the "scale separation" and the presence of the "zero mode." For vanishing TV terms, a natural scaling to satisfy $\lim_{\delta\to 0} t_i^\delta = t_i = 0$ is $|\hat{x}_j^\delta - \hat{x}_i^\delta| \propto \delta$ $(\forall j \in \partial i)$. Once this scaling is assumed, we realize that the components of the Hessian that are directly related to the clusters diverge. Let us define by $\hat{S}_\alpha$ the set of TV terms corresponding to cluster α, i.e., $\hat{S}_\alpha = \{i | (\{i\} \cup \partial i) \subset S_\alpha\}$. Hence, by construction and for all $\alpha \in C$, all components of $D_\alpha^\delta \equiv \lambda_T \left(\partial_{\boldsymbol{x}}^2 \sum_{i\in\hat{S}_\alpha} t_i^\delta\right)_{S_\alpha S_\alpha}$ are scaled as $1/\delta$ and, thus, diverge as $\delta \to 0$. The remaining terms are retained as $O(1)$. According to this "scale separation," we decompose the Hessian as $H^\delta = D^\delta + F^\delta$, where D^δ is the direct sum of the diverging components in the naively extended space; $D^\delta = \bigoplus_\alpha D_\alpha^\delta$; and F^δ consists of the remaining $O(1)$ terms.

We denote the basis of the current expression by $\{\boldsymbol{e}_i\}_{i\in S_A}$, with $(\boldsymbol{e}_i)_j = \delta_{ij}$, and move to another basis that diagonalizes $D^\delta_{S_C S_C}$. Each D_α^δ has a "zero mode," and its normalized eigenvector is given by $\boldsymbol{z}_\alpha = (z_{i\alpha})$, where $z_{i\alpha} = 1/\sqrt{|S_\alpha|}$ for $i \in S_\alpha$ and 0 otherwise, in the full space. This behavior originates from the symmetry, such that the $\{t_i^\delta\}_{i\in\hat{S}_\alpha}$ are invariant under a uniform shift in the S_α sub-space, i.e., $x_j \to x_j + \Delta$ $(\forall j \in S_\alpha)$ for $\forall \Delta \in \mathbb{R}$.

In addition, we represent the set of normalized eigenvectors of all the other modes of D_α^δ, which have eigenvalues $\lambda_{\alpha a}$ that are proportional to $1/\delta$ and positively divergent, as $\{\boldsymbol{u}_{\alpha a}\}_{a=1}^{|S_\alpha|-1}$. Then, $\{\{\{\boldsymbol{u}_{\alpha a}\}_a, \boldsymbol{z}_\alpha\}_\alpha\}$ diagonalizes $D^\delta_{S_C S_C}$ and $\{\{\{\boldsymbol{u}_{\alpha a}\}_a, \boldsymbol{z}_\alpha\}_\alpha, \{\boldsymbol{e}_i\}_{i\in S_I}\}$ constitutes an orthonormal basis of the full space. Corresponding to this variable change, we denote $\tilde{S}_Z$, $\tilde{S}_{I+Z}$, and $\tilde{S}_{C-Z}$ as the index set of variables in the space spanned by $\{\boldsymbol{z}_\alpha\}_\alpha$, $\{\{\boldsymbol{z}_\alpha\}_\alpha, \{\boldsymbol{e}_i\}_{i\in S_I}\}$, and $\{\boldsymbol{u}_{\alpha a}\}_{\alpha,a}$, respectively. In the new expression, we can rewrite $H^\delta = D^\delta + F^\delta$ as

$$H^\delta = \left(\begin{array}{c|c} D^\delta_{\tilde{S}_{C-Z}\tilde{S}_{C-Z}} & 0 \\ \hline 0 & 0 \end{array}\right) + \left(\begin{array}{c|c} F^\delta_{\tilde{S}_{C-Z}\tilde{S}_{C-Z}} & F^\delta_{\tilde{S}_{C-Z}\tilde{S}_{I+Z}} \\ \hline F^\delta_{\tilde{S}_{I+Z}\tilde{S}_{C-Z}} & F^\delta_{\tilde{S}_{I+Z}\tilde{S}_{I+Z}} \end{array}\right), \quad (13)$$

where $D^\delta_{\tilde{S}_{C-Z}\tilde{S}_{C-Z}} = \mathrm{diag}\left(\{\lambda_{\alpha a}\}_{\alpha,a}\right)$. Because of the divergence of $D^\delta_{\tilde{S}_{C-Z}\tilde{S}_{C-Z}}$, only $F^\delta_{\tilde{S}_{I+Z}\tilde{S}_{I+Z}}$ is relevant for the evaluation of $(H^\delta)^{-1}$. These considerations yield the explicit formula of χ as

$$(H^\delta)^{-1} \to \left(\begin{array}{c|c} 0 & 0 \\ \hline 0 & \left(F_{\tilde{S}_{I+Z}\tilde{S}_{I+Z}}\right)^{-1} \end{array}\right) = \chi, \quad (14)$$

where $F = \lim_{\delta\to 0} F^\delta$.

By construction, in the reduced space to $\mathrm{span}\left(\{\boldsymbol{z}_\alpha\}_\alpha, \{\boldsymbol{e}_i\}_{i\in S_I}\right)$, $F_{\tilde{S}_{I+Z}\tilde{S}_{I+Z}}$ can be expressed as

$$F_{\tilde{S}_{I+Z}\tilde{S}_{I+Z}} = \sum_{\alpha,\beta}\left(F_{\alpha\beta}\boldsymbol{z}_\alpha\boldsymbol{z}_\beta^\top + F_{\beta\alpha}\boldsymbol{z}_\beta\boldsymbol{z}_\alpha^\top\right) + \sum_\alpha\sum_{i\in S_I}\left(F_{\alpha i}\boldsymbol{z}_\alpha\boldsymbol{e}_i^\top + F_{i\alpha}\boldsymbol{e}_i\boldsymbol{z}_\alpha^\top\right) + \sum_{i,j\in S_I} F_{ij}\boldsymbol{e}_i\boldsymbol{e}_j^\top \quad (15)$$

As the non-zero components of the zero mode $\boldsymbol{z}_\alpha$ are identically given as $1/\sqrt{|S_\alpha|}$, all these coefficients can be easily expressed by the original coefficients F_{ij}, as

$$F_{\alpha\beta} = \boldsymbol{z}_\alpha^\top F \boldsymbol{z}_\beta = \frac{1}{\sqrt{|S_\alpha||S_\beta|}} \sum_{i\in S_\alpha, j\in S_\beta} F_{ij}, \quad (16a)$$

$$F_{\alpha i} = \boldsymbol{z}_\alpha^\top F \boldsymbol{e}_i = \frac{1}{\sqrt{|S_\alpha|}} \sum_{j\in S_\alpha} F_{ji}, \quad (16b)$$

and $F_{i\alpha} = F_{\alpha i}$ by the symmetry. Now, all the components are explicitly specified. The form of χ in the original basis $\{\boldsymbol{e}_i\}_{i\in S_A}$ can be accordingly assessed by moving back from the basis $\{\{\{\boldsymbol{u}_{\alpha a}\}_a, \boldsymbol{z}_\alpha\}_\alpha, \{\boldsymbol{e}_i\}_{i\in S_I}\}$, which completes the computation.

Some additional consideration of the above computation demonstrates that we can shorten some steps and obtain a more interpretable result. We introduce a $|\tilde{S}_{I+Z}| \times |\tilde{S}_{I+Z}|$ matrix $\bar{F}$ as

$$\bar{F}_{\alpha\beta} = \sqrt{|S_\alpha||S_\beta|}F_{\alpha\beta}, \ \bar{F}_{\alpha i} = \sqrt{|S_\alpha|}F_{\alpha i}, \quad (17)$$

with the remaining components being identical to those of $F_{\tilde{S}_{I+Z}\tilde{S}_{I+Z}}$, i.e., $\bar{F}_{S_I S_I} = F_{S_I S_I}$. Eqs. (16) and (17) indicate that $\bar{F}$ is simply the matrix summing the rows and columns in each cluster to a row and a column. It is natural that $\bar{F}$ has a direct connection to χ, because the locked variables in a cluster should exhibit the same response against perturbation. In fact, the response function χ in the original basis is expressed using $\bar{F}$ as

$$\chi = \sum_{i,j\in S_I} \bar{F}^{-1}_{ij} \left(\boldsymbol{e}_i \boldsymbol{e}_j^\top + \boldsymbol{e}_j \boldsymbol{e}_i^\top \right) + \sum_{\alpha,\beta} \bar{F}^{-1}_{\alpha\beta} \sum_{i\in S_\alpha} \sum_{j\in S_\beta} \boldsymbol{e}_i \boldsymbol{e}_j^\top + \sum_{\alpha} \left(\sum_{i\in S_\alpha} \sum_{j\in S_I} \bar{F}^{-1}_{\alpha j} \boldsymbol{e}_i \boldsymbol{e}_j^\top + \sum_{i\in S_I} \sum_{j\in S_\alpha} \bar{F}^{-1}_{i\alpha} \boldsymbol{e}_i \boldsymbol{e}_j^\top \right). \quad (18)$$

Eqs. (12) and (18) constitute the main result of this paper.

Application to super-resolution imaging To test the usefulness of the developed formula, let us apply the derived expression to the super-resolution reconstruction of astronomical images. A number of recent studies have demonstrated that sparse modeling is an effective means of reconstructing astronomical images obtained through radio interferometric observations [14–17]. We adopt the same problem setting as [16] and demonstrate the efficacy of our approximation formula through comparison with the literally conducted 10-fold CV result. Here, x_i denotes the ith pixel value and A is (part of) the Fourier matrix. The dataset $\boldsymbol{y}$ is generated through the linear process $\boldsymbol{y} = A\boldsymbol{x}_0 + \boldsymbol{\xi}$ where $\boldsymbol{\xi}$ is a noise vector and $\boldsymbol{x}_0$ is the simulated image, which we infer given $\boldsymbol{y}$ and A. For details on the algorithmic implementation, see the Supplemental Material [18].

In this work, we use data for simulated EHT observations based on three different astronomical images, which are available as sample data for the EHT Imaging Challenge [19]. The images are reconstructed with $N = 10000 = 100 \times 100$ pixels and with 160, 250, and 100 μas fields of view, which are identical to the original images of Datasets 1, 2, and 3 from the EHT Imaging Challenge, respectively. We test four values for each λ_{ℓ_1} and λ_T: $\lambda_{\ell_1} \in (M/2) \times \{1, 10, 100, 1000\}$ and $\lambda_T \in (M/8) \times \{1, 10, 100, 1000\}$. M is 1910, 1910, and 2786, for Datasets 1–3, respectively. Table I shows the CVE values for the three datasets, determined by the 10-fold CV and by our approximation formula for varying λ_T. λ_{ℓ_1} is fixed to the optimal value. It is clear that the approximated CVE values accord well with the 10-fold results, even on the error-bar scale, demonstrating that our approximation formula works very well. Note that the error bar for the approximation is given by the standard deviation of the M terms in Eq. (12) divided by $\sqrt{M-1}$. The actual computational time is approximately 30 h by the 10-fold CV while it is about 3 h by our approximation formula. Thus, use of the approximation formula has a clear and significant benefit.

	$8\lambda_T/M$	1	10	100	1000
Data 1	10-fold	1.101(47)	**1.090(44)**	1.091(44)	1.455(108)
	Approx.	1.103(35)	1.078(35)	**1.074(35)**	1.429(54)
Data 2	10-fold	1.368(91)	**1.260(55)**	1.286(65)	2.843(234)
	Approx.	1.236(40)	**1.159(36)**	1.211(37)	2.998(137)
Data 3	10-fold	1.026(18)	**1.020(19)**	1.020(22)	1.235(52)
	Approx.	1.029(27)	**1.015(26)**	1.017(26)	1.227(41)

TABLE I. CVE values determined by 10-fold CV and our approximation formula against λ_T. λ_{ℓ_1} is fixed to the optimal value ($2\lambda_{\ell_1}/M = 1$, coincidentally common to all cases). The number in brackets denotes the error bar to the last digits. The optimal values are bolded.

To directly observe the reconstruction quality, in Fig. 1, we show the images at all investigated parameters, the reconstructed image at the optimal λ_{ℓ_1} and λ_T, and the errors plotted against λ_T at the optimal λ_{ℓ_1}, all for Dataset 3. Again, our approximation reproduces the 10-fold result well, and the reconstructed image is a visual match of the original image. The RSS is monotonic, but the approximated LOOE is not. This implies that the LOOE factor computed through Eq. (18) appropriately reflects the penalty induced by the TV term.

Conclusion In this paper, we have developed an approximation formula for the CVE of a sparse linear regression penalized by ℓ_1 and TV terms, and demonstrated its usefulness in the reconstruction of simulated black hole images. The range of application of our formula is wide and readers are strongly encouraged to employ this formula in their own work. It is also straightforward to generalize the developed formula to other types of TV. Examples of the generalization for the so-called anisotropic TV [9] and the square TV, $T_{\rm ani} = \sum_i \sum_{j\in\partial i} |x_j - x_i|$ and $T_{\rm sq} = \sum_i \sum_{j\in\partial i} (x_j - x_i)^2$, respectively, are presented in the Supplemental Material [18].

The key concept of our formula, perturbation between the LOO and full systems, is very general and can be applied to more general statistical models and inference frameworks [20]. The development of practical formulas for those cases will facilitate higher levels of modeling and computation.

Acknowledgements We would like to express our sincere gratitude to Mareki Honma and Fumie Tazaki for their helpful discussions. We thank members of the EHT Imaging Working Group, in particular, Katie Boumann, for preparing the EHT Imaging Challenge website [19, 21]. We also thank Andrew Chael and Lindy Blackburn for writing a simulation software to produce sample data sets [22]. This work was supported by JSPS KAKENHI, Grant Numbers 26870185 (TO), 25120008 (SI), and 25120013 (YK). KA is supported by the Postdoctoral Fellowships for Research Abroad program from JSPS, and by a grant from the National Scientific Foundation (AST-1614868).

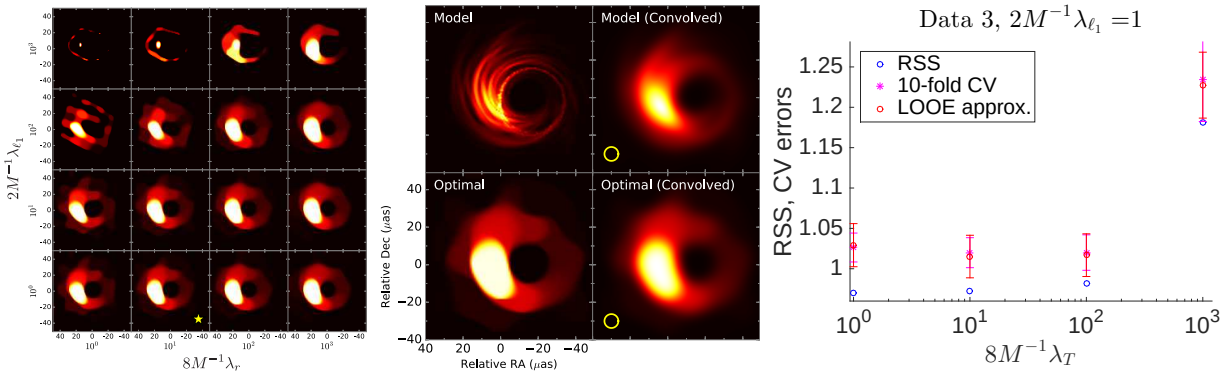

FIG. 1. Super-resolution imaging results for Dataset 3 based on model image of supermassive black hole at center of nearby elliptical galaxy, M87. (Left) Images for all investigated parameters; the star-marked panel is obtained at the optimum parameters. (Center) Original images (top) and reconstructed images (bottom) at optimal parameters ($2M^{-1}\lambda_{\ell_1} = 1, 8M^{-1}\lambda_T = 10$). The images are convolved with a circular Gaussian beam on the right-hand side, the full width at half maximum (FWHM) of which is 25% of the nominal angular resolution of the EHT and corresponds to the diameters of the yellow circles. This coincides with the optimal resolution minimizing the mean square error between them. (Right) Plot of CVEs and RSS against λ_T at $2M^{-1}\lambda_{\ell_1} = 1$. The RSS increases monotonically with growing λ_T, while the CVEs have a minimum at a central value.

* obuchi@c.titech.ac.jp

[1] I. Rish and G. Grabarnik, *Sparse Modeling: Theory, Algorithms, and Applications* (CRC Press, 2014).

[2] T. Hastie, R. Tibshirani, and M. Wainwright, *Statistical Learning with Sparsity: The Lasso and Generalizations* (CRC Press, 2015).

[3] http://sparse-modeling.jp/index_e.html

[4] J. Mairal, F. Bach, and J. Ponce, Sparse modeling for image and vision processing, arXiv:1411.3230v2.

[5] R. Tibshirani, Regression shrinkage and selection via the lasso, *J. Royal. Statist. Soc. B.*, **58**, 267 (1996).

[6] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, Least angle regression, *Ann. Stat.*, **32**, 407 (2004).

[7] L. I. Rudin, S. Osher, and E. Fatemi, Nonlinear total variation based noise removal algorithms, *Physica D* **60**, 259 (1992).

[8] A. Chambolle, An algorithm for total variation minimization and applications, *J. Math. Imaging Vision*, **20**, 89 (2004).

[9] A. Beck and M. Teboulle, Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems, *IEEE Trans. Image Process.*, **18**, 2419 (2009).

[10] T. Obuchi and Y. Kabashima, Cross validation in LASSO and its acceleration, *J. Stat. Mech.*, **2016**, 053304 (2016).

[11] http://www.eventhorizontelescope.org

[12] F. Vincent et al., High-angular-resolution and high-sensitivity science enabled by beamformed ALMA, arXiv:1309.3519.

[13] K. Akiyama et al., 230 GHz VLBI observations of M87: Event-horizon-scale structure during an enhanced very-high-energy γ-ray state in 2012, *Astrophys. J.*, **807**, 150 (2015).

[14] Y. Wiaux, L. Jacques, G. Puy, A. M. M. Scaife, and P. Vandergheynst, Compressed sensing imaging techniques for radio interferometry, *Mon. Not. R. Astron. Soc.*, **395**, 1733 (2009).

[15] F. Li, T. J. Cornwell, and F. de Hoog, The application of compressive sampling to radio astronomy I. Deconvolution, *A&A*, **528**, A31 (2011).

[16] M. Honma, K. Akiyama, M. Uemura, and S. Ikeda, Super-resolution imaging with radio interferometry using sparse modeling, *Publ. Astron. Soc. Japan*, **66**, 1 (2014).

[17] M. Honma et al., Imaging black holes with sparse modeling, *JPCS*, **699**, 012006 (2016).

[18] URL where Supplemental Material is placed.

[19] Our datasets 1, 2, and 3 correspond to the sample datasets 1, 2, and 5, respectively, available at http://vlbiimaging.csail.mit.edu/imagingchallenge

[20] Y. Kabashima, T. Obuchi, and M. Uemura, Approximate cross–validation formula for Bayesian linear regression, arXiv:1610.07733.

[21] K. Bouman et al., Computational Imaging for VLBI Image Reconstruction, The IEEE Conference on Computer Vision and Pattern Recognition, 913 (2016).

[22] A. A. Chael et al., High-resolution Linear Polarimetric Imaging for the Event Horizon Telescope, The Astrophysical Journal, **829**, 15 (2016).